

Lunar Landing Operational Risk Model

Chris Mattenberger^{a*}, Blake Putney^a, Randy Rust^b, Brian Derkowski^b

^aValador, Inc. Palo Alto, CA, USA

^bNASA Johnson Space Center, Houston, TX, USA

Abstract: Characterizing the risk of spacecraft goes beyond simply modeling equipment reliability. Some portions of the mission require complex interactions between system elements that can lead to failure without an actual hardware fault. Landing risk is currently the least characterized aspect of the Altair lunar lander and appears to result from complex temporal interactions between pilot, sensors, surface characteristics and vehicle capabilities rather than hardware failures. The Lunar Landing Operational Risk Model (LLORM) seeks to provide rapid and flexible quantitative insight into the risks driving the landing event and to gauge sensitivities of the vehicle to changes in system configuration and mission operations. The LLORM takes a Monte Carlo based approach to estimate the operational risk of the Lunar Landing Event and calculates estimates of the risk of Loss of Mission (LOM) – Abort Required and is Successful, Loss of Crew (LOC) – Vehicle Crashes or Cannot Reach Orbit, and Success. The LLORM is meant to be used during the conceptual design phase to inform decision makers transparently of the reliability impacts of design decisions, to identify areas of the design which may require additional robustness, and to aid in the development and flow-down of requirements.

Keywords: PRA, Operational Risk, Risk Informed Design, Conceptual Design Phase

1. INTRODUCTION

In order to achieve a high degree of reliability and safety in complex aerospace engineering systems, engineers must consider the risk implications of design decisions early in the conceptual phase of projects and must have an in-depth understanding of the inherent system risk [1]. Moreover, in highly mass-constrained systems, such as the Altair lunar lander, it becomes of the utmost importance to effectively use available mass to increase the reliability of the system.

As part of NASA's Constellation Program, the Altair lunar lander is to serve as the vehicle to return man to the surface of the moon. The role of the spacecraft is analogous to the Lunar Module (LM) of the Apollo program, but the performance requirements have been increased considerably. Unlike the LM, Altair is designed to be capable of transporting four astronauts to any place on the moon for a period of seven days instead of only transporting two astronauts to select locations on the moon for between one and three days [2]. These performance requirements increase the necessary complexity of the system and the effect is compounded by the constraint on the control mass of the lander determined by the selection of the Ares V as the launch vehicle within the programmatic architecture. To meet these ambitious performance requirements while staying within the control mass, the Lunar Lander Program Office (LLPO) adopted a risk informed design approach during the conceptual phase of the project. Using this approach, the team first designed a minimally functional lander which could complete the mission with high mass margins, but with an unacceptable risk of LOM or LOC during the First Lunar Design Analysis Cycle (LDAC1). Then, this approach sought to improve vehicle robustness by focusing the efforts of LDAC2 on the LOC risk drivers of the design, using the unallocated mass margin to buy-down risk as indicated by the team's probabilistic risk assessment (PRA) of the vehicle.

Standard PRAs usually employed to verify that a product meets requirements are too resource intensive and too slow to keep up with the speed at which the design is maturing in the conceptual phase; while classical qualitative methods do not provide the level of detail and granularity required by the designers to make high-quality risk informed decisions. Moreover, often times PRAs will over emphasize the contribution of hardware failures to the overall risk profile. Examination of the root causes of historical spacecraft failures indicates that only a small portion of all system failures can be attributed to random part failures [19,20,21]. To perform the risk analysis necessary to implement risk informed design during the conceptual phase of the project, the LLPO utilized the Valador Reliability Tool (VRT) [3]. Centered on interaction between risk analyst and designer, the VRT uses a component based method to identify all initiating events which may lead to a LOM/LOC event. The analytic tool is able to quickly produce results at the component, system, and vehicle level while easily adapting to the rapidly changing design space through increased designer – risk analyst interaction [4]. For each component, the tool takes in a predicted failure rate, expected duty cycle, and applicable system responses to a loss of component functionality. In addition, a top-down approach is used to identify operational and event risks which apply to the entire system, but do not stem from a hardware failure. Examples of this include: Micro-Meteoroid (MM) and Orbital Debris, Software, Low Lunar Orbit (LLO) Docking, and Landing. These failure modes are added directly to the model.

During LDAC1, these risks were identified and assigned a conservative best-guess placeholder value that would be revisited and refined as the design progressed. As part of the design efforts for LDAC2, the lunar orbit and surface affects of MM and LLO Docking risks were quantified through analysis and updated. As the system matured and became more robust, the Landing risk began to emerge as the number one risk driver of the vehicle for both LOM and LOC. As the risk picture is left incomplete with the primary risk driver modeled as a conservative estimate, this motivated the development of the LLORM for the purpose of gaining a better understanding of the intrinsic nature of this event risk. The LLORM seeks to determine whether this estimate is a risk driver due to the conservatism of the estimate or if it is a risk driver due to an actual design deficiency. It will aid in identifying areas of the design which may be overly conservative and will help to identify potential risk mitigation strategies.

2. MODEL DEVELOPMENT

The development of the LLORM began with the purpose to gain insight into the necessary vehicle processes and the interactions between pilot, vehicle and surface during a lunar landing attempt. The development of this model was also influenced by the nature of the design project. As the LLORM was being developed for use during the conceptual phase of the design project, it became an important guiding principal to develop a tool that could flexibly respond to the rapidly changing design and concept of operations. Moreover, the conceptual nature of the design directed the development of the tool such that it could aid in focusing and structuring technical decisions, rather than certifying the design meets some risk requirement. To achieve this, the development was focused on a generic lunar lander as opposed to a high-fidelity model tailored specifically to Altair. In addition, the development of the model framework attempted to isolate the processes into distinct modules, much like in object oriented programming [5,6], which could be easily and iteratively ‘upgraded’ to higher fidelity models as guided by the emerging operational risk.

The development plan called for an iterative process which would add fidelity to the model, guided by the inherent risk of the processes, in progressive phases of development. To begin Phase 0, a historical review was performed which examined the Apollo experience drawing from various source documents [7,8], works of literature examining lunar landing attempts [9], Apollo mission reports [10,11,12,13,14,15] and interviews with area experts. The goal of Phase 1 was to establish a generic operational framework with placeholder process modules based upon the historical review. The next step, during Phase 2, is to replace the placeholder modules with actual performance models for a specific configuration which would yield a meaningful quantitative result. During Phase 3, the capability of the model is to be increased such that modules can take on a range of configurations to allow for the exploration of the sensitivities of the LOM/LOC scoring to various input parameters.

Ultimately, the LLORM will be capable of performing Monte Carlo style simulations of the total process required to land on the moon while giving the design team the capability to explore the risk-design space and make well informed design decisions.


Currently, Phase 0 and Phase 1 have been completed and Phase 2 is underway. The following sections explore in further detail the work completed and describe the plan forward for the remaining Phases.

2.1. Historical Review (Phase 0)

The precursor investigation focused mainly on the experience of the Apollo program, drawing from actual flight experience, historical design documents, and post-programmatic analyses. In addition, the results of the preliminary system hazard analysis were examined to determine possible failure modes during the landing event. The goal of this activity was to allow the inherent risk of the process to guide the development of the model.

This allowed for the identification of the major processes required to successfully land on the moon. After undock from the crew capsule, the lander undergoes a series of checkouts to confirm the health and state of the lander while still in a Stable Parking Orbit. Next, the lander fires the descent main engine to perform the De-orbit Maneuver, leaving the relative safety of the stable parking orbit and coasts until the Powered Descent Initiation occurs to begin the Braking Phase. At some point near the surface, the lander performs a Pitch-Up Maneuver and begins the Approach Phase. Now, the processes begin to occur in much faster succession, starting with a Sensor Scan of the area to determine the location of possible Safe Landing Sites. Next, there is some degree of Human Interaction with the lander to observe the sensor data, orient the data relative to the current situation, decide upon a landing point to designate, and physically act to implement this decision. The lander must now respond to this input and Translate to the designated landing site. As the lander approaches the landing site, one last final check is performed to confirm the safety of the site and once confirmed the lander will then proceed to the Terminal Descent Phase and land on the surface of the moon. A graphical representation of this general process is shown in Figure 1.

Figure 1. Descent & Landing Graphical Representation



Evidence gathered during this phase also confirmed the initial supposition that the risk drivers of the landing event were not hard failures of the components needed to complete the landing event. This idea is reinforced when considering the probability of a random failure being modeled by an exponential type failure probability based upon a failure rate and time as seen in Equation 1.

$$P(\text{Failure}) = 1 - e^{-\lambda * t} \quad (1)$$

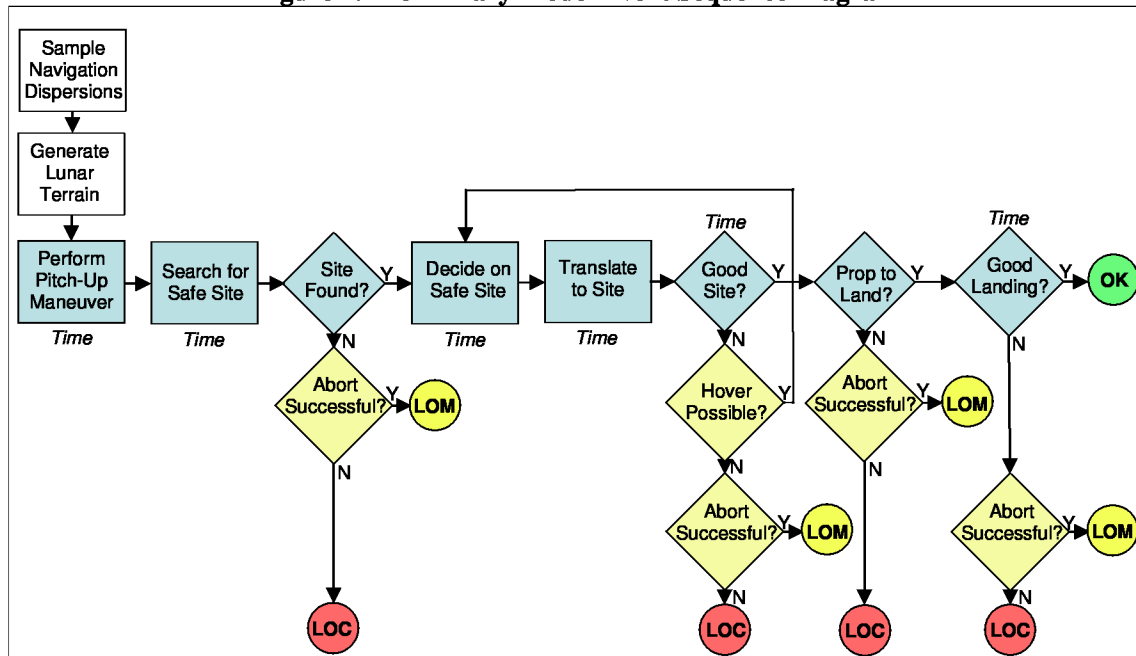
The probability of a random failure during the landing event is very small due to the narrow window of time each of the specific components is needed to function without failure, beginning after the

planned component checkouts prior to the start of the landing event. In addition, it is reasonable to assume that a generic lunar lander would also have some degree of robustness or fault tolerance for components critical to the success of landing.

Moreover, the review yielded the insight that a key parameter involved in this event is propellant remaining or if the engine is assumed to fire at a nearly constant throttle during the approach phase, then the metric of importance is Time. The amount of time the engine has been firing can be thought to indicate how much fuel is remaining as well as the position of the lander along the nominal trajectory. The landing attempt will be terminated when the propellant runs low.

Taking these ideas into consideration, the review led to the formation of the basic operational flow as seen in the event sequence diagram in Figure 2. Those processes which accrue time are noted.

Figure 2: Preliminary Model Event Sequence Diagram



The next step in the development of the risk model was to decide on a style of risk modeling which would most easily handle the inherent risk of the system as well as the envisioned iterative approach. In addition to the insights gleaned from the historical review, the conceptual nature of the design influenced this decision as well. An event tree model was initially considered for this process. Branches in the event tree would be defined by success criteria (i.e. Pilot chooses the landing point in time T); however, the success criteria of the branches in the landing process would have a high degree of interdependence due to the dynamic nature of this mission phase. This interdependence would lead to an explosion in the complexity of the model and would render the model impractical as a design tool. This led to the natural selection of a Monte Carlo style approach to the risk modeling. Such an approach allows the landing event to be broken down into step-by-step process modules which can be easily altered or enhanced as the design and model evolve as more about the event is understood. Moreover, this style of approach, when compared to a bottom up fault tree oriented model, forces the risk analyst to focus on establishing the success criteria of events and thus forces the analyst to interact intensely with the designers rather than focusing on developing the failure rates of components, which can often be a more turn-key actuarial exercise for a developing design.

Additionally, as the resources for the conceptual design phase are at a minimum, it was important to take an approach which would be able to leverage the work already being done by other members of

the team. A key point of emphasis considering this is that much of these processes were concurrently being examined by other parts of the design team; however, their approach to these processes was from the perspective of design and performance optimization. The goal here is to leverage this existing work by examining the work from a risk perspective, focusing on off-nominal behavior and incorporating the necessary distilled elements of the existing models into the LLORM.

The Monte Carlo approach allows for an iterative process in which capabilities can be built up over time as directed by the preliminary results of the model. Now, the operational framework, based upon the historical review, can be established. Next, the existing outside models can be integrated into the risk model in a very basic fashion. Then, guided by the current results, the fidelity of specific process modules can be augmented. And ultimately, the completed model will be able to provide comprehensive results for a range of possible input parameters.

2.2. Operational Framework (Phase 1)

Implementation of the model takes place within the Extend v6 [16] process simulator software and Microsoft Excel [17]. Extend controls the overall Monte Carlo process as well as each of the individual realizations and sorts and counts the outcome of each trial. Excel contains the logic and most of the process modules. During an individual realization, Extend first generates a random number which is then passes to Excel. This number is used to seed the random number function in Excel used in various modules as well as to control the input / output interface between Extend and Excel. Once the Visual Basic for Applications code has finished running, Excel then passes a number back to Extend. Based upon the value returned, Extend sorts and counts the results of the trial.

The following sections describe the existing basic operational framework of the LLORM. Each of these processes is implemented in the model as a stand alone module and captures the risk identified during the historical review as well as risks discovered during the implementation of the model.

2.2.1 Navigational Dispersions

This first module is meant to capture the risk associated with successfully arriving at the correct location to begin the Pitch-Up Maneuver. Uncertainties in engine burn performance, errors in the knowledge of the gravitational field of the moon, map-tie error, guidance system noise, and accelerations below the detectable limits of the inertial measurement unit among others all contribute to the potential positional dispersions of the lander. This module seeks to quantify these types of errors that are introduced beginning at the De-orbit Maneuver, through the Powered Descent Initiation phase, and result in a probabilistic distribution of possible lander locations at the time of the Pitch-Up Maneuver. This module also takes into account methods of reducing potential navigational dispersions. Most notably, the ability of the vehicle to engage in Terrain Relative Navigation (TRN) will greatly reduce the amount of uncertainty in this area.

Currently, the model neglects the effects of navigational dispersions given that the TRN system will be sufficient to overcome the sources of error. This assumes that the components required for successful function of the TRN system are working properly. This is based upon navigational dispersion analysis conducted by an Altair subsystem design team which showed the potential dispersions for the Altair lunar lander with a functioning TRN system to be on the order of several meters. This is insignificant when compared to the overall size of the Landing Field which is on the order of hundreds of meters.

2.2.2 Terrain Generation

This module incorporates a probabilistically accurate lunar surface to serve as a random Landing Field with possible Safe Landing Sites determined by a combination of surface location characteristics of slopes and hazards (rocks and craters) and vehicle hazard/slope tolerances as well as the vehicle landing gear footprint. In addition, the process is influenced by the results of the Navigational

Dispersions and the nominal mission profile. This module randomly determines how many Safe Landing Sites are present in the Landing Field and randomly places these sites in the Landing Field.

The initial version of the model used a simple distribution lookup function. A random number is generated based upon a uniform distribution between one and zero. This value is then used to determine a corresponding number of Safe Landing Sites to be placed in the Landing Field. Section 2.3 Outside Model Integration describes in detail how this initial placeholder function has been enhanced through the incorporation of a distillation of an outside high-fidelity model.

2.2.3 Pitch-Up Maneuver

This module captures the uncertainty in the amount of time required to perform the Pitch-Up Maneuver and to place the vehicle in the correct state to begin the sensor scan of the Landing Field. The longer this process takes, the less time will be available to complete other time-critical operations. Further, a lengthening of this process alters the position of the vehicle along the nominal trajectory when subsequent position sensitive operations occur.

The current module in the tool is a placeholder uniform distribution which randomly selects an amount of time between 15s and 30s.

2.2.4 Sensor Performance

When the Pitch-Up Maneuver is complete, the Hazard Detection Sensor determines what locations the system 'believes' to be possible Safe Landing Sites. This module captures the uncertainty in the ability of the sensor to accurately detect slopes and hazards as well as uncertainties in the amount of time required to perform a sensor sweep, process the data, and display the results to the crew. If no safe sites are found, then the system must perform a nominal abort and attempt to return to orbit within the given constraints on time to abort and propellant remaining.

This module models dispersions in both time and successful operation. The amount of time required to complete this process is determined by starting with the best case operation time and then adding a random delay based upon sampling a uniform distribution between zero and the best case operation time. The best case time was a placeholder guess and Section 2.3 Outside Model Integration discusses in further detail how the fidelity of this time uncertainty modeling has been increased.

To determine the successful operation of the hazard sensor in the model, this module examines every possible Landing Site within the Landing Field. Depending on the true suitability of the site, safe or bad, the module chooses the appropriate conditional probability to sample using a uniform random number between zero and one. The key influential properties of this process are the frequency of false positives ($P(+|-)$), a hazard is detected, but does not exist and false negatives ($P(-|+)$), a site is thought to be safe, but actually contains a landing hazard. Currently, the module is using a set of placeholder conditional probabilities for the chance of the hazard detection sensor to function as follows: $P(-|-) = 0.90$ and $P(+|-) = 0.10$, $P(-|+) = 0.001$ and $P(+|+) = 0.999$.

2.2.5 Human Interaction

The uncertainty parameter of this module focuses on the amount of time required for the crew to confirm/designate a Landing Point. This process can be decomposed into four main activities following an OODA (Observe, Orient, Decide, Act) loop [18]. First, the crew must observe the output from the processing of the Hazard Detection Sensor data. The crew must then orient themselves properly to this data and understand the current situation. Next, the crew must decide upon a course of action and finally must physically act upon the decision and either confirm the nominal Landing Point or re-designate to an identified Safe Landing Site.

The current module contains a placeholder uniform distribution between 10s and 30s.

2.2.6 Hazard Relative Navigation

This module determines the capability of the system to shift or re-designate the Landing Point if the nominal landing point is not a safe landing site. The capability is a function of the current time, or position, along the nominal trajectory as well as the amount of Delta-V margin, or additional propellant, which has been included in the vehicle design. This module works in concert with the Human Interaction module to limit the possible Safe Landing Sites to only those which can be reached within the propellant budget at that time. The module will also output how much time is required to reach the landing site. In addition, this module functions in the off-nominal case when a Landing Site has been reached, but fails to pass the Final Landing Site Check and the lander must translate to a new landing site if the lander is capable of zeroing out all descent velocity and hovering.

At present, this module is an extremely simplified placeholder process that is not well defined and allows the vehicle to reach any landing site within the landing area at anytime. The time to reach the landing site is accrued as a function of distance from the nominal landing point to the re-designated landing site. Potential improvement ideas for this module are described in Section 2.5 Future Work.

2.2.7 Final Landing Site Check

This event occurs just prior to the initiation of the Terminal Descent Phase and serves as the final confirmation of a safe site prior to committing to a landing attempt. While, the exact description of this 'check' is left undefined, the intent is that some sort of sensor, video camera, or even a human eyeball will allow the crew to verify that the site is indeed safe and clear of unacceptable landing hazards or slopes. If the landing site is determined to be unacceptable, then depending on the capabilities of the lander and propellant remaining, either a nominal abort is performed or the lander enters a hover mode and must attempt to translate to the nearest safe site. In addition, a sufficient amount of propellant must remain to make a landing attempt; otherwise the vehicle must perform a nominal abort and attempt to return to orbit.

This is modeled through the use of placeholder conditional probabilities. The greatest impacts to the overall mission success are false positives ($P(+|-)$), incorrectly identifying a bad site as safe, and false negatives ($P(-|+)$), incorrectly identifying a safe site as bad. Currently, the module uses the following conditional probabilities: $P(+|-) = 0.001$ and $P(-|-) = 0.999$, $P(-|+) = 0.001$ and $P(+|+) = 0.999$.

2.2.8 Landing Parameter Dispersions

This performance dispersion captures the uncertainty surrounding the ability of the vehicle to successfully complete the Terminal Descent Phase and to land the vehicle on the surface within the acceptable landing parameters determined by the capabilities and structure of the vehicle. These parameters include rates, velocities and orientation of the vehicle at touchdown. If this process fails, then an emergency abort must be performed as a final attempt to save the crew.

Currently, this module contains a best guess probability of successful touchdown with $P(S) = 0.999$ based upon expert judgment.

2.2.9 Abort Modes

Failure to complete certain processes or running out of propellant to perform a safe landing or allow for safe abort will cause the system to perform an abort and attempt to return to orbit. Within the model, there is a 'nominal' abort and an 'emergency' abort corresponding to the severity of the initiating failure. Primarily, the nominal abort case occurs if no safe landing site is found or the vehicle has run out of propellant to perform a safe landing or allow for a safe abort. It is used in

situations where the abort mode can be anticipated and the need for an abort is not highly time sensitive. The emergency abort occurs for failure modes which are highly time critical and the preparation time for an abort is small. Examples of this include critical component failures as well as a failure of the system to complete the Touchdown Event within appropriate landing parameters.

The best guess probabilities of success for the nominal and off-nominal abort cases are, respectively, $P(S) = 0.999$ and $P(S) = 0.9$ based upon expert judgment.

2.3. Outside Model Integration (Phase 2)

With the completion of the first iteration of the model, attention can now be focused on the integration of uncertainty analysis already being performed by different elements of the design team as well as incorporation of high fidelity performance models being developed outside of the design team. Each of these independent analyses has its own boundary conditions, assumptions, and dependencies. Some may assume the best case or nominal case which can lead to masking of actual risk drivers. Others may make worst case assumptions which can lead to an over designed system if these analyses are simply stacked upon one another. During this phase, attention must be given to such considerations in order to effectively integrate and leverage existing work while minimizing the resources employed to increase the fidelity of the modules as guided by the emerging risk of the design.

In order to determine where to focus the efforts of model development, the Phase 1 Operational Framework placeholder model was examined in detail. This analysis focused on the relative frequencies of the various failure modes encountered by the model after running approximately 10,000 realizations of the Monte Carlo simulation. In addition, the sensitivities of the output of the model to changes in module distributions were studied to gauge the relative impact that each module has upon a successful outcome. This highlighted the processes that have the greatest impact on the success of the landing event, indicating which process modules should be improved upon first. Additionally, consideration is given to those modules which are obviously at a significantly lower level of fidelity than the rest of the current model. Furthermore, due to the limited resources of any conceptual design phase, further model development can often be dictated by the availability of the designers to work in concert with the risk analyst to understand, distill, and integrate outside models.

Another important consideration is to what degree outside analyses should be incorporated. Keeping with the guiding principals, it is desirable to start with the simplest possible integration method and only increase complexity if warranted. The level at which an outside model is integrated into the LLORM is based upon the degree of dependence upon boundary conditions or the results of other processes. If the outside model is completely independent of other operations, then it can be included as a single probability or distribution. If the outside model has a slight degree of dependence on other operations, then it could be included as a set of distributions or a lookup table. More complicated interactions between processes may require a limited capability model to be integrated as a function. It may even be the case that to fully capture the dependencies the complete version of the model should be included within the LLORM.

Much of the current state of the model is still made up of placeholder processes; however, some work has been done to incorporate outside models in several key processes. A simple example of this considers the uncertainty in the amount time required to complete the hazard detection as described in Section 2.2.4 Sensor Performance. To increase the fidelity of this placeholder, the risk analyst interviewed a design engineer currently engaged in research and development efforts of an innovative flash LIDAR (Light Detection and Ranging) sensor which is being designed with use on a lunar lander in mind. Real world experience from actual flight tests of this sensor yielded a set of surrogate data which were used to come up with the best case operating time. These data are based upon the time recorded for the sensor and data processing system to collect information about a field of rocks and slopes in the desert, then process this information and display it on a monitor.

A more in-depth example of this type risk analyst – designer interaction and model integration was alluded to in Section 2.2.2 Terrain Generation. In this example, the work of the Autonomous Landing and Hazard Avoidance Technology (ALHAT) group at the Jet Propulsion Laboratory was examined, distilled, and incorporated into the LLORM. The ALHAT team had already developed a model to create probabilistically accurate lunar surface maps based upon lunar terrain types for slopes and hazards (Smooth Mare, Rough Mare, Hummocky Uplands, Rough Uplands). The ALHAT team can then determine what area on this map would be safe for a lunar lander based upon key vehicle parameters (Vehicle Footprint Dispersion Ellipse, Rock Height Tolerance, and Slope Tolerance).

The output of the model focused on predicting the probability of the existence of at least one possible safe landing site in a random realization of this process, while the needs of the LLORM require a distribution of the number of safe landing sites present in a field. By utilizing the raw output of the ALHAT terrain model, the risk analyst is able to calculate a rough estimate of the number of safe site in each of the 500 trials of a specific terrain type and vehicle class provided by the ALHAT team. A histogram can then be created from these data which can be easily used to create a probabilistic distribution of the number safe sites. This can be sampled using a uniform distribution between zero and one. The implementation of this distribution lookup function can be seen in Figure 3.

Figure 3. Integration Example: Terrain Generation Distribution Lookup Function

```
ALHAT = Rnd(seed)
'Case 20
If ALHAT <= 0.008 Then
    numSites = 0
ElseIf ALHAT <= 0.022 Then
    numSites = 1
ElseIf ALHAT <= 0.074 Then
    numSites = 2
ElseIf ALHAT <= 0.168 Then
    numSites = 3
ElseIf ALHAT <= 0.308 Then
    numSites = 4
ElseIf ALHAT <= 0.502 Then
    numSites = 5
ElseIf ALHAT <= 0.726 Then
    numSites = 6
ElseIf ALHAT <= 0.882 Then
    numSites = 7
ElseIf ALHAT <= 0.98 Then
    numSites = 8
Else
    numSites = 9
End If
```

In order to produce meaningful results, much more effort must be dedicated to the integration of existing uncertainty and performance models to upgrade many of the placeholders in the current model. The general plan for the continuation of efforts similar to those described here can be found in Section 2.5 Future Work.

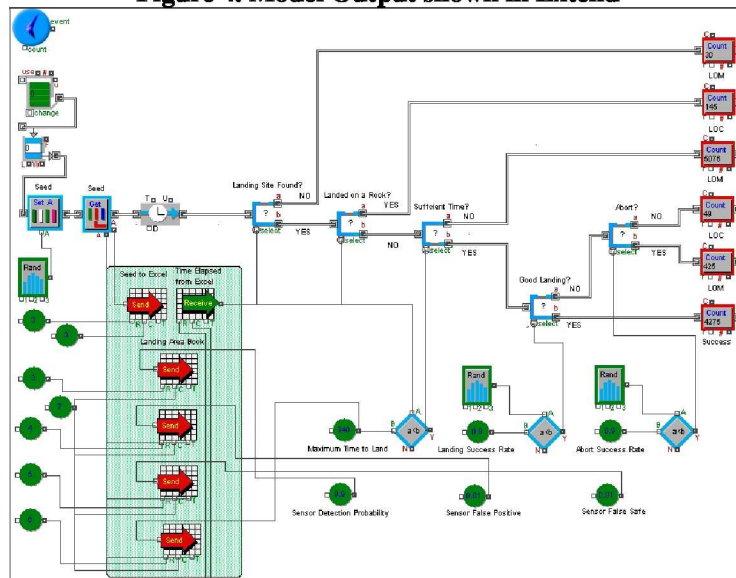
2.4. Sample Results (Phase 3)

This section goes into greater detail about the envisioned outputs of the LLORM and the ways the LLORM can be used to aid in the design process of the Altair lunar lander. The point LOM/LOC estimate for a particular design are calculated and displayed in the Extend portion of the model as seen below in Figure 4. To calculate an estimate of the risk of LOM/LOC, the model simply counts the number of realizations resulting in each outcome and divides the total by the total number of realizations. Determining how many total realizations to run is important to the validity of each estimate as it is important for the probabilities to converge while allowing access to the results in a short time frame. Results for the placeholder model were obtained using 10,000 realizations which demonstrated some degree of convergence and allowed for rare events to manifest.

This portion of the model also shows details about the initiating failure which resulted in a LOM or LOC. Understanding the breakdown of the initiating failures can yield significant insights into the landing process. Such results can be used to determine which areas of the design may be deficient and

warrants further analysis. Examining the initiating failure mode distribution can also be a useful exercise to challenge the intuition of the design team and to serve as a spring board for thoughtful technical discussion. It is envisioned that once the model has been completed the results will help to increase the confidence level in the initial conservative estimate of the risk of the Landing Event.

Figure 4. Model Output shown in Extend



A more substantial and influential use of the tool will be to perform sensitivities studies of the estimates for LOM/LOC to changes in the design of the vehicle. This will help the design team to focus on ‘the differences that make a difference’ by highlighting areas of the vehicle that have the most impact on the overall success of the mission. The value of this exercise goes beyond a simple quantitative estimate and focuses more deeply upon the relative differences between design solutions.


Identifying key parameters will help to focus future trade studies and analyses. Examination of these parameters can yield significant insights into the inherent nature of the system risk. For instance, if the LOM/LOC scoring appears to be insensitive to changes in a specific module or process, then perhaps mass should not be used to increase the robustness of this process or this could identify an area of the design which may be overly conservative. Sensitivity studies can also inform the designer as to which processes would most greatly benefit from additional performance margin, how mission design can impact the likelihood of success, or could influence the requirement setting for future components of the design. For example, a sensitivity study could be performed to determine how sensor false positive performance changes the LOM/LOC picture. Knowing this information can allow the designers to set very specific requirements for the operation of this sensor and in turn affect the development of this technology or how it is used during a nominal mission. Overall, these types of studies will have more of an impact than the point estimate for LOM/LOC because they compare design options relatively and focus the key parameters that substantially affect the integrated system performance and risk.

2.5. Future Work

Specific plans have been made to work with various members of the design team as well to work with researchers outside the scope of Altair. Increased model fidelity can be achieved by extracting insights from experts in the field of sensor performance, human factors, trajectory analysis, and others. Examples of this interaction include working further with the ALHAT team to determine the probabilistic performance parameters of the hazard detection sensor. Can the sensor operation be distilled into a conditional probability or is this an over-simplification? To what degree is the sensor performance dependent upon the terrain class of the landing field or position along the nominal

trajectory? Another example of future work includes taking notional ideas and filling in the important details extracted from rigorous analysis. Figure 5 shows a notional concept of a function to be used in the Hazard Relative Navigation module which would take in time and additional propellant available and produce an output of an area in which landing point re-designation is possible as well as the time required to reach a chosen landing point. A similar performance analysis has already been conducted, but now it must be understood, distilled, and integrated from a risk perspective while aligning assumptions and boundary conditions with the other integrated uncertainty analyses.

Figure 5. Notional Concept for Hazard Relative Navigation



The next major milestone in the development of this model will be to gain buy-in from the subsystem designers, project decision makers, and the PRA community through demonstration of a LLORM concept model. The concept model will not use any placeholder modules, but will have a limited functionality and a limited range of potential sensitivity parameters. The goal here is to demonstrate the validity of the approach and examine the model results in a narrow case.

3. CONCLUSION

Ultimately, the LLORM is meant to serve as a flexible design tool which will allow for the rapid exploration of the design space with minimal impacts upon project resources by leveraging analyses already under way. The approach is geared towards integrating 'stove-piped' aspects of the design and bringing together members the design team across subsystem boundaries to encourage focused technical interchanges and to yield system level insights about the inherent system risk. To achieve these goals, the fidelity of each of the key processes must be increased far beyond that of a placeholder and as a whole the fidelity must be balanced across each module.

This paper has described the form, function, and intended use of a process oriented operational risk model. The model will focus on 'the differences that make a difference' and exploration of the risk implications of design decisions, rather precise prediction of the overall risk of the Landing Event. In addition, the paper detailed an innovative approach to the development of such a model which is fast enough to keep up with the speed of the evolving design during the conceptual phase of the project, yet maintains a sufficient level of fidelity through the incorporation of existing performance analyses. At a later phase in the project, a high fidelity, detail oriented, integrated landing simulation will be warranted to verify the risk-performance of the vehicle, but at this early point in the project life, it would prove to be too cumbersome and too resource intensive.

Risk informed design helps the design team by aiding in the efficient and effective allocation of project, mission and campaign resources. By focusing the efforts of the design team on key reliability metrics, it is possible to inform the intuition of designers, to yield insights into the inherent nature of the risk of key processes, and to develop solutions balancing risk, performance, and cost. These types of PRA tools are not meant to make design decisions, but rather to enlighten the designers of the risk

implications of their decisions, to spawn thoughtful discussion and to develop concrete rationale for design decisions. And what's more, this style and approach can be applied to a wide range of design projects outside the scope of manned space exploration.

As the political realities which have forced NASA to make crew safety and mission reliability a design driving requirement begin to spread to other industries, the demands upon complex systems to operate reliably and to continue operating safely in light of a failure will continue to increase. These demands will further require reliability analysis to evolve, innovate, and mature tools which can effectively, transparently and rapidly add value to design projects earlier and earlier in the conceptual design phase of complex engineering projects.

Acknowledgements

The authors would like to gratefully acknowledge the support and leadership of the Lunar Lander Program Office at Johnson Space Center. The authors would also like to thank their fellow analysts at Valador and colleagues at the Jet Propulsion Laboratory and Johnson Space Center for all of their hard work and efforts put into this design project. Particular acknowledgement is given to the ALHAT team for aiding in outside model development and integration.

References

- [1] J. Miller, J. Leggett, and J. Kramer-White, "*Design Development Test and Evaluation Considerations for Safe and Reliable Human Rated Spacecraft Systems*", NASA, 2008, Hampton, VA.
- [2] T. Jones. "*Shooting for the Moon*", Aerospace America, May 2008, pp. 20-22, (2008).
- [3] B.F. Putney, E. Tavernetti, J.R. Fragola, and E. Gold. "*Reliability Tool for a Preliminary Quantified Functional Risk and Hazard Analysis*", Proceedings of the Reliability and Maintainability Symposium, 2009.
- [4] C.J. Mattenberger. "*Risk Informed Design Process & Design Team – Analyst Interaction*", Proceedings of the Reliability and Maintainability Symposium, 2010.
- [5] F. Friedman and E. Koffman, "*Problem Solving, Abstraction, and Design Using C++*", Pearson Education, 2004, Boston, MA.
- [6] M. Feldman and E. Koffman, "*Ada 95: Problem Solving and Program Design*", Addison Wesley Longman, 1999, Reading, MA.
- [7] D. Cheatham and F. Bennett, "*Apollo Lunar Module Landing Strategy*", NASA.
- [8] W. Malloy and W. Everett, "*Apollo Operations Handbook: Lunar Module*", Grumman Aerospace Corporation, 1969, Bethpage, NY.
- [9] D. Mindell, "*Digital Apollo: Human and Machine in Spaceflight*", The MIT Press, 2008, Cambridge, MA.
- [10] G. Low, "*Apollo 11 Mission Report*" NASA, 1969, Houston, TX.
- [11] J. McDivitt, "*Apollo 12 Mission Report*" NASA, 1970, Houston, TX.
- [12] J. McDivitt, "*Apollo 14 Mission Report*" NASA, 1971, Houston, TX.
- [13] J. McDivitt, "*Apollo 15 Mission Report*" NASA, 1971, Houston, TX.
- [14] O. Morris, "*Apollo 16 Mission Report*" NASA, 1972, Houston, TX.
- [15] O. Morris, "*Apollo 17 Mission Report*" NASA, 1973, Houston, TX.
- [16] B. Diamond, "*Extend v6: Developer's Reference*", Imagine That, 2002, San Jose, CA.
- [17] M. Kofler, "*Definitive Guide to Excel VBA*", Springer-Verlag, 2003, New York, NY.
- [18] J. Boyd, "*The Essence of Winning and Losing*", Boyd, 1995.
- [19] H. Hecht and M. Hecht, "*Reliability Prediction for Spacecraft*", Rome Air Development Center, 1985, Rome, NY.
- [20] W.F. Tosney and A. H. Quintero, "*Orbital Experience from an Integration and Test Perspective*", Journal of IEST, November/December 1998, (1998).
- [21] J. Bullman, G. Langford, and M. Benik, "*OSP-ELV Human Flight Safety Certification Study Report*", NASA, 2004, Huntsville, AL.